

Set Covering Problems with General Objective Functions

Jean Cardinal, Christophe Dumeunier

Université Libre de Bruxelles (ULB)
Computer Science Department, CP 212
B-1050 Brussels, Belgium
{jcardin, cdumeuni}@ulb.ac.be

Abstract. We introduce a parameterized version of set cover that generalizes several previously studied problems. Given a ground set V and a collection of subsets S_i of V , a feasible solution is a partition of V such that each subset of the partition is included in one of the S_i . The problem involves maximizing the mean subset size of the partition, where the mean is the generalized mean of parameter p , taken over the elements. For $p = -1$, the problem is equivalent to the classical minimum set cover problem. For $p = 0$, it is equivalent to the minimum entropy set cover problem, introduced by Halperin and Karp. For $p = 1$, the problem includes the maximum-edge clique partition problem as a special case. We prove that the greedy algorithm simultaneously approximates the problem within a factor of $(p + 1)^{\frac{1}{p}}$ for any $p \in \mathbb{R}^+$, and that this is the best possible unless $P = NP$. These results both generalize and simplify previous results for special cases. We also consider the corresponding graph coloring problem, and prove several tractability and inapproximability results. Finally, we consider a further generalization of the set cover problem in which we aim at minimizing the sum of some concave function of the part sizes. As an application, we derive an approximation ratio for a Rent-or-Buy set cover problem.

1 Introduction

The greedy strategy is one of the simplest and most well-known heuristic, which can be applied to many combinatorial optimization problems. In the case of the minimum set cover problem, it involves iteratively choosing a subset that covers a maximum number of uncovered elements. We study this algorithm on a natural family of set covering problems in which the value of a subset depends on the number of elements it covers, and a parameter p encodes the way in which these values are combined. This parameter interpolates between different versions of the set covering problem, in particular between the classical minimum set cover problem, the minimum entropy set cover problem, and the simpler problem of finding a subset of maximum size.

Intuitively, the greedy algorithm should perform better for objective functions in which more importance is given to subsets covering many elements. We give a formal support to this intuition by showing that the greedy algorithm provides a constant factor approximation for all positive values of the parameter p . We further show that this is the best we can achieve unless $P = NP$.

We first define some notations. Let V be an n -element ground set and $\mathcal{S} = \{S_1, \dots, S_k\}$ a collection of k subsets of V , whose union is V . In the minimum set cover problem, we seek a minimum size subset $\mathcal{T} \subseteq \mathcal{S}$ such that $\bigcup_{S_i \in \mathcal{T}} S_i = V$. We define a *cover* as an assignment $\varphi : V \mapsto \mathcal{S}$ of each element of V to a set of \mathcal{S} such that $v \in \varphi(v)$ for all $v \in V$. This definition allows us to define alternative objective functions for the set cover problem. Given a cover φ , let us define a *part* as a set $\varphi^{-1}(S_i)$ for some $S_i \in \mathcal{S}$. We use the following two notations: $c_i := |\varphi^{-1}(S_i)|$ is the *part size* of the i th subset S_i with respect to φ , and $a_v := |\varphi^{-1}(\varphi(v))|$ is the size of the part containing the element v , with $v \in V$.

We define a new family of set cover problems in which we aim at maximizing the *mean* $M(\{a_v : v \in V\})$ of the values a_v . There exist many definitions of the mean $M(\{a_1, a_2, \dots, a_n\})$ of a set of numbers.

The most widely used definition is the *arithmetic mean*: $M_1(\{a_1, a_2, \dots, a_n\}) := \frac{1}{n} \sum_{i=1}^n a_i$. Another well-known definition is the *geometric mean*: $M_0(\{a_1, a_2, \dots, a_n\}) := (a_1 \cdot a_2 \cdot \dots \cdot a_n)^{\frac{1}{n}}$. Finally, we also consider the *harmonic mean*: $M_{-1}(\{a_1, a_2, \dots, a_n\}) := n / (\sum_{i=1}^n a_i^{-1})$. The arithmetic, geometric, and harmonic means are special cases of the *generalized mean*:

$$M_p(\{a_1, a_2, \dots, a_n\}) = \left(\frac{1}{n} \sum_{v \in V} a_v^p \right)^{\frac{1}{p}} = \left(\frac{1}{n} \sum_{i: c_i \neq 0} c_i^{p+1} \right)^{\frac{1}{p}}. \quad (1)$$

This value is the arithmetic mean for $p = 1$, and the harmonic mean for $p = -1$. It is well-known that the limit of the generalized mean for $p \rightarrow 0$ is equal to the geometric mean. The generalized mean with parameter p is also called the normalized L_p -norm¹.

Definition 1 (Maximum p -mean set cover). Given an n -element ground set V and a collection $\mathcal{S} = \{S_1, \dots, S_k\}$ of subsets of V whose union is V , find a cover $\varphi : V \mapsto \mathcal{S}$ that maximizes $M_p(\{a_v : v \in V\})$, where $a_v := |\varphi^{-1}(\varphi(v))|$, and M_p is the generalized mean of parameter p .

Special Cases

Interestingly, letting $p = -1$ (harmonic mean) or $p = 0$ (geometric mean) yields set cover problems that are already known: the harmonic mean version is the minimum set cover problem, while the geometric mean version is the *minimum entropy set cover* problem [6]. A special case of the maximum p -mean set cover problem for $p = 1$ has recently been introduced in the form of a graph coloring problem [9].

Minimum Set Cover. The maximum harmonic mean set cover problem can be cast as $\min_{\varphi} \sum_{v \in V} \frac{1}{a_v}$. We can rewrite this objective function as $\sum_{v \in V} \frac{1}{a_v} = \sum_{S_i \in \mathcal{S}} \sum_{v \in \varphi^{-1}(S_i)} \frac{1}{c_i} = |\{S_i : c_i \neq 0\}|$. Hence the maximum harmonic mean set cover problem is the standard minimum set cover problem.

This problem is among the most studied NP-hard problems. It has long been known to be approximable within a factor $H_{\max_i |S_i|}$ with the greedy algorithm. The first proof is from Johnson [20]. Lovász [23] obtained the same factor with a different method. Later, Chvátal extended the result to the weighted set cover problem [8], in which the subsets S_i have nonuniform costs. A number of papers show that the logarithmic approximation guarantee is likely to be optimal. Lund and Yannakakis [24] first proved that the problem is not approximable within $\log n/4$ unless $\text{NP} \subseteq \text{DTIME}(n^{\text{polylog}(n)})$. This result has been improved to $(1 - o(1)) \ln n$ by Feige [10], under the hypothesis $\text{NP} \not\subseteq \text{DTIME}(n^{O(\log \log n)})$. Raz and Safra [27], and Alon, Moshkovitz, and Safra [1] proved inapproximability results for factors of the form $c \ln n$ for some constant c under the hypothesis $\text{P} \neq \text{NP}$. These results are consequences of new PCP characterizations of NP.

Minimum Entropy Set Cover. Let us now consider the geometric mean version: $\max_{\varphi} (\prod_{v \in V} a_v)^{\frac{1}{n}}$. We relate this mean to the *entropy* of the discrete probability distribution found by dividing each part size by n :

$$\begin{aligned} - \sum_{i=1}^k \frac{c_i}{n} \log \frac{c_i}{n} &= - \sum_{v \in V} \frac{1}{n} \log \frac{a_v}{n} \\ &= \log n - \frac{1}{n} \sum_{v \in V} \log a_v \\ &= \log n - \log M_0(\{a_v : v \in V\}). \end{aligned}$$

¹ We use the word p -mean here, in order to avoid confusion with the “minimum L_p -norm set cover” problem [15].

Thus the maximum geometric mean set cover problem is equivalent to the problem of minimizing the entropy of the partition. This problem is known as the *minimum entropy set cover* problem. It has been introduced by Halperin and Karp [19], and has applications in the field of computational biology. They proved that the problem was approximable within a constant additive term with the greedy algorithm. Improving on this work, Cardinal, Fiorini, and Joret [6] provided a simple analysis showing that the constant was at most $\log_2 e \simeq 1.4427$ bits, and that this was the smallest additive error achievable in polynomial time, unless $P = NP$. The minimum entropy vertex cover [7] and minimum entropy graph coloring [5] problems, which are special cases of minimum entropy set cover, have been studied by the same authors.

Maximum-Edge Clique Partition. In a recent publication [9], Dessmark, Jansson, Lingas, Lundell, and Persson studied the *maximum-edge clique partition* (Max-ECP) problem. In this problem, we aim to partition a graph G into cliques in order to maximize the number of edges whose endpoints are in the same clique of the partition. This is an implicit set cover problem, in which the subsets S_i are the cliques of G , and the function to maximize is:

$$\sum_{i=1}^k \binom{c_i}{2} = \frac{1}{2} \left(-n + \sum_{i=1}^k c_i^2 \right) = \frac{n}{2} (M_1(\{a_v : v \in V\}) - 1).$$

Thus the problem can be seen as an implicit maximum p -mean set cover problem for $p = 1$. They show that the problem is 2-approximable on perfect graphs using the greedy algorithm, and that it is not approximable within a factor $n^{1-O(1/(\log n)^\gamma)}$ for some constant γ in polynomial time unless $NP \subseteq ZPTIME(2^{(\log n)^{O(1)}})$.

Max-Max and Max-Min Set Cover. When $p \rightarrow \infty$, the maximum p -mean set cover problem involves finding a cover in which the largest part has maximum size. This is a trivial problem, unless the subsets in \mathcal{S} are not given explicitly, like in the graph coloring problem. For $p \rightarrow -\infty$, the problem is that of maximizing the size of the smallest part, thus solving $\max_{\varphi} \min_{v \in V} a_v = \max_{\varphi} \min_{i: c_i \neq 0} c_i$. This problem seems much more challenging. We will refer to it as the max-min set cover problem.

Our results

We show in section 2 that for any $p \in \mathbb{R}^+$, the maximum p -mean set cover problem is approximable within a factor of $(p + 1)^{1/p}$. This factor is less than e for all positive values of p , hence this can be seen as a *robust e -approximation* for all p -means with positive p . This result generalizes the approximability results of Cardinal et al. [6] for the case $p \rightarrow 0$, and of Dessmark et al. [9] for $p = 1$. We also prove that this is the best we can achieve in polynomial time unless $P = NP$, using a powerful reduction due to Feige et al. [10,11]. When p is negative, we show that the performance of the greedy algorithm degrades. We give an inapproximability result for max-min set cover.

Graph coloring problems can be seen as implicit set cover problems in which the subsets S_i are the maximal independent sets of the graph. The subsets are not given explicitly, which would cause an exponential blowup in the problem size, but rather implicitly, from the graph structure. We define the maximum p -mean graph coloring problem in this natural way. Special cases of the maximum p -mean graph coloring problem include the standard minimum coloring problem ($p = -1$), the minimum entropy coloring problem [5] ($p \rightarrow 0$), the maximum-edge clique partition problem [9] ($p = 1$), and the maximum independent set problem ($p \rightarrow +\infty$). In Section 3 we give approximability and inapproximability results for this problem.

The maximum p -mean set cover problem involves maximizing the sum of the $(p + 1)$ th powers of the part sizes, as can be seen in equation (1). In section 4, we consider weighted instances, and a further

generalization of the set cover problem, in which we aim at minimizing the sum of some concave function of the part sizes. We give a closed form of the approximation ratio achieved by the greedy algorithm for this general class of problems, and apply this result to the case of the Rent-or-Buy set cover problem [12].

Related works

Minimum sum set cover. In the minimum sum set cover problem we aim to find an ordering of the subsets that minimizes the average *cover time* of an element of the ground set, where the cover time of an element is the index of the first subset covering it. This problem was first considered in its graph coloring version [4]. Feige, Lovász, and Tetali [11] gave an elegant proof of the fact that greedy is a 4-approximation algorithm, and that this was the best one could hope for unless $P = NP$. They also studied the related minimum sum vertex cover problem, for which they provided a 2-approximation algorithm.

Generalizations of minimum sum set cover. Munagala, Babu, Motwani, and Widom [25] introduced the *pipelined set cover* problem. In this problem, we aim to find an ordering of the subsets in \mathcal{S} that minimizes the L_p -norm of the vector (R_i) , where R_i is the number of elements that are not contained in any of the first $(i - 1)$ subsets. For $p = 1$, the problem is equivalent to the minimum sum set cover problem. They generalize the technique of Feige et al. [11] to prove a $4^{\frac{1}{p}}$ -approximation.

More recently, Golovin, Gupta, Kumar, and Tangwongsan [15] considered another minimum L_p -norm set cover problem. This variant involves finding an ordering of the subsets minimizing the L_p -norm of the cover time vector. This problem is a simultaneous generalization of the minimum set cover problem and the minimum sum set cover problem. They prove that the greedy algorithm provides a $O(p)$ -approximate solution, and that this is the best possible, up to a constant factor, unless $NP \subseteq DTIME(n^{O(\log \log n)})$.

Graph Coloring. The greedy algorithm for set cover translates to the MaxIS algorithm for graph coloring, in which a maximum independent set is iteratively chosen as new color class. This algorithm has in particular been analyzed for the minimum sum [4] and minimum entropy [5,6] graph coloring problems.

Recently, Fukunaga, Halldórsson, and Nagamochi [13] initiated the study of a very general family of minimum cost graph coloring problems, similar to what we propose in section 4. They proved that any minimum cost graph coloring problem in this family is 4-approximable on weighted interval graphs, provided that the cost function is both monotone and concave. The proposed algorithm iteratively removes a maximum i -colorable subgraph, where i is doubled at each iteration.

In another recent contribution, Fukunaga, Halldórsson, and Nagamochi [12] introduced the *Rent-or-Buy* coloring problem in vertex-weighted graphs, in which the cost of a color class is the minimum between 1 and the total weight of the class. This models situations in which each color class has to be paid for either by “buying” it for a fixed cost, or “renting” it for a price proportional to its size. They gave, among other results, a 2-approximation for this problem in perfect graphs. We consider the set cover version of this problem in section 4.

Clique Partitioning with Value-polymatroidal Costs. Gijswijt, Jost, and Queyranne [14] recently studied clique partitioning problems with *value-polymatroidal* cost functions. A function f over the subsets of V is said to be value-polymatroidal whenever $f(\emptyset) = 0$, f is non-decreasing and for every subsets S and T with $f(S) \geq f(T)$, and every u in $V \setminus (T \cup S)$, the inequality $f(S + u) - f(S) \leq f(T + u) - f(T)$ holds. They define the cost of a clique partition as the sum of the cost of each clique. They prove, among other results, that this problem is solvable in polynomial time on interval graphs.

Minimum L_p -norm problems. Azar, Epstein, Richter, and Woeginger [2] studied approximation algorithms for a scheduling problem in which we aim to *minimize* the L_p -norm of the part sizes. A similar problem has been studied by Azar and Taub [3], who proposed all-norm approximation algorithms. Although similar in spirit, the goal is different than ours, since we instead seek the most “nonuniform” distribution, with *maximum* L_p -norm.

A number of other problems with general cost functions have been studied, such as facility location [17]. Due to space constraints, we do not give more details here.

2 Approximability

Lemma 1. *The maximum p -mean set cover problem for $p \in \mathbb{R}$ is approximable in polynomial time within a factor of*

$$\left(\frac{n^{p+1}}{\sum_{j=1}^n j^p} \right)^{\frac{1}{p}}. \quad (2)$$

Proof. We consider an optimal cover φ_{OPT} , and a part $C_i = \varphi_{\text{OPT}}^{-1}(S_i)$ in this cover, of size $|C_i| = c_i$. We define $a'_v := |\varphi^{-1}(\varphi(v))|$ for the cover φ returned by the greedy algorithm.

We first suppose that $p \geq 0$, and give a lower bound on the value of the cover φ restricted to C_i . We do so by examining the elements of C_i in the order in which they are covered by the greedy algorithm, breaking ties arbitrarily. The first covered element $v_1 \in C_i$ must belong to a part of size at least c_i in φ , since C_i can be chosen as a new part, and the greedy algorithm chooses the largest part. Hence $a'_{v_1} \geq c_i$. Similarly, the second element v_2 of C_i that is covered by greedy must belong to a class of size at least $c_i - 1$. Hence $a'_{v_2} \geq c_i - 1$. In general, for the k th element v_k covered by the greedy algorithm, $a'_{v_k} \geq c_i - k + 1$. Thus we have

$$\sum_{v \in C_i} (a'_v)^p \geq \sum_{j=1}^{c_i} j^p. \quad (3)$$

Letting $a_v := |\varphi_{\text{OPT}}^{-1}(\varphi_{\text{OPT}}(v))|$, the corresponding value for φ_{OPT} is $\sum_{v \in C_i} a_v^p = c_i^{p+1}$, hence we get the following upper bound

$$\frac{\sum_{v \in C_i} a_v^p}{\sum_{v \in C_i} (a'_v)^p} \leq \frac{c_i^{p+1}}{\sum_{j=1}^{c_i} j^p}. \quad (4)$$

This ratio is increasing with c_i , and holds for all the parts C_i of φ_{OPT} . Letting $c_i = n$ and taking the p th root gives the result.

A similar reasoning holds for $p < 0$, with the direction of inequalities (3) and (4) reversed. \square

The approximation ratios for various values of p and n are given in Fig. 1. We next give a constant upper bound on the approximation ratio in the case $p \geq 0$. We need the following lemma.

Lemma 2. *For $p \in \mathbb{R}^+$ and $n \in \mathbb{N}$,*

$$\sum_{j=1}^n j^p \geq \frac{n^{p+1}}{p+1}.$$




Fig. 1. Approximation ratios for the greedy algorithm.

Proof. The inequality holds for $p = 0$. For $p > 0$, it can be checked graphically that approximating the sum by an integral yields a lower bound:

$$\sum_{j=1}^n j^p = \sum_{j=0}^n j^p > \int_0^n x^p dx = \frac{n^{p+1}}{p+1}.$$

□

Combining lemmas 1 and 2 proves the following theorem. Tightness can be proved using known tight examples for special cases (see for instance [6]).

Theorem 1. *The maximum p -mean set cover problem is approximable in polynomial time within a factor of $(p+1)^{\frac{1}{p}}$ for $p \in \mathbb{R}^+$. This bound is asymptotically tight.*

Note that $\lim_{p \rightarrow +\infty} (p+1)^{\frac{1}{p}} = 1$, hence in the case of $p \rightarrow +\infty$, the approximation ratio is equal to 1. This formalizes the trivial observation that if our goal is to maximize the size of the largest part, then the greedy algorithm returns an optimal solution. Also, $\lim_{p \rightarrow 0} (p+1)^{\frac{1}{p}} = e$, which proves that the greedy algorithm approximates the minimum entropy set cover within an additive term of $\log e$ bits. This was shown by Cardinal, Fiorini, and Joret [6]. Finally, for $p = 1$, the greedy algorithm returns a 2-approximation. A proof of this result was given by Dessmark, Jansson, Lingas, Lundell, and Persson [9].

We now turn to the case $p < 0$. We know that the greedy algorithm approximates the problem for $p = -1$ within a logarithmic factor. The following result shows that the performance of greedy degrades dramatically as p becomes smaller.

Theorem 2. *The maximum p -mean set cover problem is approximable in polynomial time within a factor of $n^{1-\frac{1}{q}} \zeta(q)^{\frac{1}{q}}$ for any real $p = -q < -1$, where $\zeta(q) = \sum_{j=1}^{\infty} j^{-q}$ is the Riemann zeta function.*

Proof. We consider expression (2) in lemma 1 and replace p by $-q$:

$$\left(\frac{n^{1-q}}{\sum_{j=1}^n j^{-q}} \right)^{-\frac{1}{q}} = \left(\frac{\sum_{j=1}^n j^{-q}}{n^{1-q}} \right)^{\frac{1}{q}} \leq \left(\frac{\zeta(q)}{n^{1-q}} \right)^{\frac{1}{q}} = n^{1-\frac{1}{q}} \zeta(q)^{\frac{1}{q}}. \quad (5)$$

□

The bound is asymptotically tight if we replace n by $\max_i |S_i|$. Note that we need $q > 1$, otherwise the Dirichlet series defining the zeta function does not converge. In particular, when $q = 1$ (and thus $p = -1$), we have the harmonic series, which is the approximation ratio for the minimum set cover problem.

An interesting special case is when $p = -2$. This means that the cost of a part of size c_i in the cover is $1/c_i$. In that case, the approximation ratio of the greedy algorithm becomes

$$n^{1-\frac{1}{2}} \zeta(2)^{\frac{1}{2}} = \pi \sqrt{\frac{n}{6}}. \quad (6)$$

We now show that the approximability result in theorem 1 for positive values of p is the best we can hope for, unless $P = NP$. We need the following lemma, which is a simple consequence of the convexity of the function $f(x) = x^{p+1}$.

Consider two sorted sequences $c_1 \geq c_2 \geq \dots \geq c_k$ and $c'_1 \geq c'_2 \geq \dots \geq c'_k$. We say that (c_i) *dominates* (c'_i) if

$$\sum_{i=1}^j c_i \geq \sum_{i=1}^j c'_i \quad \forall j \in \{1, 2, \dots, k\}. \quad (7)$$

Lemma 3. *If (c_i) dominates (c'_i) , then for any $p \in \mathbb{R}^+$,*

$$\sum_{i=1}^k c_i^{p+1} \geq \sum_{i=1}^k (c'_i)^{p+1}. \quad (8)$$

Theorem 3. *It is NP-hard to approximate the maximum p -mean set cover problem within a factor less than $(p+1)^{\frac{1}{p}}$ for $p \in \mathbb{R}^+$.*

Proof. Feige, Lovasz, and Tetali [11] gave a procedure for transforming a 3SAT-6 formula into a set system (V, \mathcal{S}) with the following properties:

- each subset $S_i \in \mathcal{S}$ has size n/t for a certain parameter t ,
- if the formula is satisfiable, then there exists an exact cover of V with t subsets,
- if the formula is δ -satisfiable, that is, if at most a fraction δ of the clauses can be satisfied, then every i subsets of \mathcal{S} cover at most a fraction $(1 - (1 - 1/t)^i) - \varepsilon$ of the elements of V , for $i \in \{1, 2, \dots, at\}$ and any choice of the constants $\varepsilon > 0$ and $a > 0$.

Given a formula known to be either satisfiable or δ -satisfiable, the problem of distinguishing between the two is NP-hard [11]. Using the transformation above, we show that a polynomial algorithm with an approximation ratio less than $(p+1)^{\frac{1}{p}}$ for maximum p -mean set cover would solve this problem.

If the formula is satisfiable, then V can be covered by exactly t disjoint sets of \mathcal{S} . From Lemma 3, this is the optimal solution. The part sizes c_i in this solution satisfy

$$\sum_{i=1}^k \left(\frac{c_i}{n}\right)^{p+1} = \sum_{i=1}^t \left(\frac{1}{t}\right)^{p+1} = \frac{1}{t^p}. \quad (9)$$

We now suppose the formula is only δ -satisfiable. We consider the distribution in which the i th part covers a fraction

$$\left(1 - (1 - 1/t)^i\right) - \left(1 - (1 - 1/t)^{i-1}\right) = \frac{1}{t} \left(1 - \frac{1}{t}\right)^{i-1}$$

of the elements of V , for $i \in \{1, 2, \dots, at\}$, and the remaining parts cover exactly a fraction $\frac{1}{t} \left(1 - \frac{1}{t}\right)^{at}$ each. We denote by r the number of remaining parts, so that the sum of the fractions equals 1. From Lemma 3 and the properties of the reduction, this distribution dominates all other achievable distributions. Therefore the following upper bound holds.

$$\sum_{i=1}^k \left(\frac{c_i}{n}\right)^{p+1} \leq \sum_{i=0}^{at} \left(\frac{1}{t} \left(1 - \frac{1}{t}\right)^i\right)^{p+1} + r \left(\frac{1}{t} \left(1 - \frac{1}{t}\right)^{at}\right)^{p+1} \quad (10)$$

$$\simeq \frac{1}{t^{p+1}} \sum_{i=0}^{at} e^{-\frac{(p+1)i}{t}} + \frac{r}{t^{p+1}} e^{-a(p+1)}. \quad (11)$$

We can approximate the sum by an integral :

$$\sum_{i=0}^{at} e^{-\frac{(p+1)i}{t}} \simeq \int_0^{at} e^{-\frac{(p+1)x}{t}} \cdot dx = \frac{t}{p+1} \left(1 - e^{-a(p+1)}\right). \quad (12)$$

The value r is the number of parts of size $\frac{1}{t} \left(1 - \frac{1}{t}\right)^{at} \simeq \frac{1}{t} e^{-a}$ needed to cover a fraction $1 - \sum_{i=0}^{at} \frac{1}{t} \left(1 - \frac{1}{t}\right)^i \simeq e^{-a}$ of the elements. Thus $r \sim t$, and

$$\frac{r}{t^{p+1}} e^{-a(p+1)} \simeq \frac{1}{t^p} e^{-a(p+1)}.$$

Note that since the constant t can be assumed to be arbitrary large [11], the approximations above are arbitrarily accurate. Hence expression (11) can be made arbitrarily close to:

$$\frac{1}{t^p} \left(\frac{1}{p+1} \cdot \left(1 - e^{-a(p+1)}\right) + e^{-a(p+1)} \right). \quad (13)$$

Now by choosing a sufficiently large, the ratio between (13) and (9) can be made arbitrary close to $p+1$. The gap between the p -means is obtained by taking the p th root. \square

In the case $p \rightarrow 0$, the above inapproximability proof shows that the additive $\log e$ error term is best possible (unless $P = NP$) for the minimum entropy set cover problem. This was also shown previously by Cardinal, Fiorini, and Joret [6].

Although we do not have a precise inapproximability threshold for negative values of p , we can prove the following result for $p \rightarrow -\infty$. That is the max-min set cover problem, in which we aim to maximize the size of the smallest part.

Theorem 4. *It is NP-hard to approximate the max-min set cover problem within any constant factor.*

Proof. The proof uses the same reduction as the proof of theorem 3. We consider set systems (V, \mathcal{S}) constructed from a 3SAT-6 formula, such that there exists an exact cover with t parts of size $\frac{n}{t}$ if the formula is satisfiable, and every i subsets of \mathcal{S} cover at most a fraction $(1 - (1 - 1/t)^i) - \varepsilon$ of the elements, for $i \in \{1, 2, \dots, at\}$, if the formula is δ -satisfiable. But this means that in the latter case, at least at subsets are needed to cover V . This implies that there is a part of size at most $\frac{n}{at}$. Since a can be chosen arbitrarily greater than any constant, the gap can be made arbitrarily large. \square

3 Graph Coloring

We now define the graph coloring variant of the maximum p -mean set cover problem.

Definition 2 (Maximum p -mean graph coloring). *Given a simple, undirected graph $G = (V, E)$, find an assignment $\varphi : V \mapsto \mathbb{N}$ of colors to vertices such that adjacent vertices receive different colors, and $M_p(\{a_v : v \in V\})$ is maximized, where $a_v := |\varphi^{-1}(\varphi(v))|$ and M_p is the generalized mean with parameter p .*

The greedy algorithm extends naturally to what is referred to as the MaxIS algorithm, in which a maximum independent set is iteratively removed from the graph. This procedure can run in polynomial time only if at each step we can find a maximum independent set in polynomial time. This is true for large families of graphs, such as perfect graphs [16], and claw-free graphs [26]. We thus have the following corollary of theorem 1 (the proof of tightness is omitted).

Corollary 1. *The maximum p -mean graph coloring problem restricted to perfect or claw-free graphs is approximable in polynomial time within a factor of $(p+1)^{\frac{1}{p}}$ for $p \in \mathbb{R}^+$. This bound is asymptotically tight.*

It may happen that we only have an approximate algorithm for the maximum independent set problem. Then the following result applies. Proofs are given in appendix A.

Theorem 5. *If the maximum independent set problem can be approximated within a factor ρ in polynomial time, then the maximum p -mean graph coloring problem is approximable within a factor of $\rho(p+1)^{\frac{1}{p}}$ in polynomial time.*

Corollary 2. *The minimum entropy coloring problem [5] is approximable in polynomial time within an additive error of $\log_2(\Delta+2) - 0.14226$ on graphs with maximum degree Δ .*

In the max-min graph coloring problem, that is when $p \rightarrow -\infty$, we aim to maximize the size of the smallest color class. Using a recent polynomial algorithm from Kierstead and Kostochka to construct equitable $\Delta+1$ -colorings [22], we have the following approximability result.

Corollary 3. *The max-min graph coloring problem can be approximated in polynomial time within a factor $(1 + O(\frac{1}{n}))^{\frac{\Delta+1}{\chi}}$ on graphs of order n , maximum degree Δ , and chromatic number χ .*

The maximum independent set problem is the special case of minimum p -mean coloring in which $p \rightarrow +\infty$. It is therefore not surprising that the general coloring problem is not well approximable for any positive value of p , as the following lemma shows.

Lemma 4. *If the maximum independent problem set cannot be approximated in polynomial time within $n^{1-\varepsilon}$ for some $\varepsilon = \varepsilon(n)$, then the maximum p -mean graph coloring problem with $p \in \mathbb{R}^+$ cannot be approximated in polynomial time within $n^{1-(2+\frac{1}{p})\varepsilon}$.*

Proof. If the maximum independent set cannot be approximated within $n^{1-\varepsilon}$, then we can safely assume that this holds for graphs having an independent set of size $\alpha \geq n^{1-\varepsilon}$. In such a graph, we consider the coloring obtained with a $n^{1-t\varepsilon}$ -approximation algorithm for maximum p -mean coloring, for some constant t to be fixed later.

The optimal solution in this graph has value at least $(\alpha^{p+1})^{\frac{1}{p}}$. Thus the value A of the coloring satisfies

$$A \geq \frac{(\alpha^{p+1})^{\frac{1}{p}}}{n^{1-t\varepsilon}}. \quad (14)$$

We now consider the largest color class in this coloring, and denote its size by h . We then get the following upper bound on A :

$$A \leq \left(\frac{n}{h} h^{p+1}\right)^{\frac{1}{p}} = h n^{\frac{1}{p}}. \quad (15)$$

Putting this together, we obtain

$$h n^{\frac{1}{p}} \geq \frac{(\alpha^{p+1})^{\frac{1}{p}}}{n^{1-t\varepsilon}} \geq \frac{(n^{(1-\varepsilon)(p+1)})^{\frac{1}{p}}}{n^{1-t\varepsilon}} \quad (16)$$

$$h \geq n^{\left(t-1-\frac{1}{p}\right)\varepsilon}. \quad (17)$$

Letting $t = 2 + \frac{1}{p}$, we obtain an independent set of size at least n^ε , which is a $n^{1-\varepsilon}$ -approximation for the maximum independent set problem, a contradiction. \square

Applying this lemma and using a result from Khot [21], we obtain the following.

Theorem 6. *The maximum p -mean graph coloring problem, for $p \in \mathbb{R}^+$, is not approximable in polynomial time within a factor $n^{1-O(1/(\log n)^\gamma)}$ for some constant γ unless $NP \subseteq ZPTIME(2^{(\log n)^{O(1)}})$.*

A similar result for $p \rightarrow 0$ was proved by Cardinal et al. [5]. The special case $p = 1$ was proved by Dessmark et al. [9].

We end our discussion of the graph coloring problems with the equivalent problem in the complement of the graph G , which we call the *maximum p -mean clique partition* problem. The Max-ECP problem corresponds to the special case $p = 1$. Gijswijt, Jost, and Queyranne [14] provided a $O(n^3)$ dynamic programming algorithm for finding a partition of interval graphs in cliques that minimizes the sum of a value-polymatroidal cost. Unfortunately, our objective function do not fall in that class, since the equivalent minimization problem involves minimizing a concave decreasing cost function, and value-polymatroidal functions must be non-decreasing. However, the correctness of their dynamic programming solely relies on the fact that an optimal partition always contain a maximal clique. This is true in our case as well, at least for $p > 0$, and is a consequence of lemma 3. Thus the algorithm can be applied and we get the following results.

Theorem 7. *The maximum p -mean clique partition problem with $p \in \mathbb{R}^+$ can be solved in $O(n^3)$ time on interval graphs.*

Corollary 4. *The Max-ECP problem [9] can be solved in $O(n^3)$ time on interval graphs.*

4 Further Generalizations

Weighted variant. We first observe that theorems 1 and 3 also hold for a weighted version of the minimum p -mean set cover problem. In this problem, the elements of v have a weight $w(v)$. The objective function is the same, except that a_v is now defined as $w(\varphi^{-1}(\varphi(v)))$. We can observe that the approximability proofs

above still hold using a simple reduction for integer weights. Given a weighted instance, we can transform it into an unweighted instance by replacing each element $v \in V$ by $w(v)$ copies of it, each belonging to the same subsets as v . Then each copy of the duplicated elements must belong to the same part of the (greedy or optimal) solution. Otherwise, from lemma 3, some elements can be reassigned so that the p -mean increases. The argument extends to rational and, by continuity, real weights.

General costs. Following the definition of Fukunaga, Halldórsson, and Nagamochi [13] for minimum cost colorings, we now consider a much more general family of set cover problems. In these problems, we aim to minimize a sum of some concave function $f(c_i)$ of the part sizes. The functions f are concave in the sense that they are discrete restrictions of concave functions $f : \mathbb{R}^+ \mapsto \mathbb{R}$. We also assume $f(0) = 0$. Setting $f(c_i) = -c_i^{p+1}$, for instance, yields a problem similar to the maximum p -mean set cover problem, without the $1/p$ exponent. The definition of this new family is as follows.

Definition 3 (Set cover with general costs). *Given an n -element ground set V and a collection $\mathcal{S} = \{S_1, \dots, S_k\}$ of subsets of V whose union is V , find a cover $\varphi : V \mapsto \mathcal{S}$ that minimizes $\sum_{i=1}^k f(c_i)$, where $c_i := |\varphi^{-1}(S_i)|$ and f is a concave function.*

Concavity implies that we seek a distribution of the part sizes that is as unbalanced as possible. In particular, the following generalization of lemma 3 holds.

Lemma 5. *Given two nonincreasing sequences (c_i) and (c'_i) , such that (c_i) dominates (c'_i) , and a concave function f , we have $\sum_{i=1}^k f(c_i) \leq \sum_{i=1}^k f(c'_i)$.*

Although the approximation ratio obtained with the greedy algorithm depends on the function f , we can give a simple expression of it.

Theorem 8. *The set cover problem with general costs can be approximated in polynomial time within a factor of*

$$\max \left\{ \frac{1}{f(c)} \sum_{j=1}^c \frac{f(j)}{j} : 1 \leq c \leq \max_i |S_i| \right\}.$$

Proof. (sketch) Given a solution φ , we associate to each element $v \in V$ the cost $\frac{f(a_v)}{a_v}$, where $a_v = |\varphi^{-1}(\varphi(v))|$ as before. The cost of this solution is the sum $\sum_{v \in V} \frac{f(a_v)}{a_v}$. Using concavity, we can bound this sum in a greedy solution as in the proof of lemma 1: we show that the sum over the elements in a part of size c in the optimal solution is at most $\sum_{j=1}^c f(j)/j$. The ratio follows. \square

Note that we retrieve the approximation ratio H_n of minimum set cover by setting $f(c) = 1$ if $c > 0$ and $f(0) = 0$. This result also encompasses our analyses of the approximability of minimum entropy and maximum p -mean set cover.

We now give an application of this result to a new problem. In this problem, we suppose that the cost of assigning an element of V to a subset S_i is 1 if S_i covers a lot of elements, but is proportional to its size if the fraction of elements covered by S_i is small. More precisely, if the fraction c_i/n of elements covered by S_i is greater than some constant $\beta < 1$, then the incurred cost is $\frac{c_i}{n}/\beta$. Otherwise, the cost is 1. Thus β defines a breakpoint, above which it is less costly to “buy” the subset than “rent” it. Hence we define the *Rent-or-Buy* set cover problem as the set cover problem with the following cost function:

$$f(c) = \begin{cases} c/(\beta n) & \text{if } c \leq \beta n, \\ 1 & \text{otherwise.} \end{cases}$$

This models situations in which for instance jobs are assigned to machines, and machines can be either bought or rented. The model was introduced recently by Fukunaga, Halldórsson, and Nagamochi as a graph coloring problem [12]. The original description of the Rent-or-Buy model was on a weighted graph, and the coloring problem was to find a coloring minimizing the sum of the values $\min\{1, w(C_i)\}$ over all color classes C_i , where $w(C_i)$ is the sum of the weights of the vertices in C_i . From our reduction of weighted instances described above, this is equivalent to our problem with $\beta = \frac{1}{w(V)}$.

Corollary 5. *The Rent-or-Buy set cover problem is approximable in polynomial time within a factor of $1 - \ln \beta$.*

Proof. We let $t = \beta n$. Let us first suppose that $c \leq t$. Then we have

$$\frac{1}{f(c)} \sum_{j=1}^c \frac{f(j)}{j} = \frac{t}{c} \sum_{j=1}^c \frac{1}{t} = 1. \quad (18)$$

Otherwise, if $c > t$, we have

$$\frac{1}{f(c)} \sum_{j=1}^c \frac{f(j)}{j} = 1 + \sum_{j=t+1}^c \frac{1}{j} = 1 + H_c - H_t \leq 1 + H_n - H_t \leq 1 - \ln \beta. \quad (19)$$

Hence from Theorem 8, this is the worst-case approximation ratio achieved by the greedy algorithm. \square

Since the greedy algorithm can be implemented to run in polynomial time on perfect or claw-free graphs, we obtain the following result on the Rent-or-Buy graph coloring problem.

Corollary 6. *The Rent-or-Buy coloring problem is approximable in polynomial time within a factor of $1 + \ln w(V)$ on perfect or claw-free graphs.*

This improves on the 2-approximation algorithm [12] when the overall weight $w(V)$ does not exceed e .

References

1. N. Alon, D. Moshkovitz, and M. Safra. Algorithmic construction of sets for k -restrictions. *ACM Transactions on Algorithms*, 2(2):153–177, 2006.
2. Y. Azar, L. Epstein, Y. Richter, and G. J. Woeginger. All-norm approximation algorithms. In *Scandinavian Workshop on Algorithms Theory (SWAT)*, volume 2368 of *Lecture Notes in Computer Science*, pages 288–297. Springer-Verlag, 2002.
3. Y. Azar and S. Taub. All-norm approximation for scheduling on identical machines. In *Scandinavian Workshop on Algorithms Theory (SWAT)*, volume 3111 of *Lecture Notes in Computer Science*, pages 298–310. Springer-Verlag, 2004.
4. A. Bar-Noy, M. Bellare, M. M. Halldórsson, H. Shachnai, and T. Tamir. On chromatic sums and distributed resource allocation. *Information and Computation*, 140(2):183–202, 1998.
5. J. Cardinal, S. Fiorini, and G. Joret. Minimum entropy coloring. In *16th International Symposium on Algorithms and Computation (ISAAC)*, volume 3827 of *Lecture Notes in Computer Science*, pages 819–828. Springer-Verlag, 2005.
6. J. Cardinal, S. Fiorini, and G. Joret. Tight results on minimum entropy set cover. In *9th. International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*, volume 4110 of *Lecture Notes in Computer Science*, pages 61–69. Springer-Verlag, 2006. To appear in *Algorithmica*.
7. J. Cardinal, S. Fiorini, and G. Joret. Minimum entropy orientations. preprint arXiv:0802.1237v1 [cs.DS], 2008.
8. V. Chvátal. A greedy heuristic for the set-covering problem. *Mathematics of Operations Research*, 4(3):233–235, 1979.
9. A. Dessmark, J. Jansson, A. Lingas, E.-M. Lundell, and M. Persson. On the approximability of maximum and minimum edge clique partition problems. *International Journal of Foundations of Computer Science*, 18(2):217–226, 2007.
10. U. Feige. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM*, 45(4):634–652, 1998.
11. U. Feige, L. Lovász, and P. Tetali. Approximating min sum set cover. *Algorithmica*, 40(4):219–234, 2004.

12. T. Fukunaga, M. M. Halldórsson, and H. Nagamochi. “rent-or-buy” scheduling and cost coloring problems. In *Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*, volume 4855 of *Lecture Notes in Computer Science*, pages 84–95. Springer-Verlag, 2007.
13. T. Fukunaga, M. M. Halldórsson, and H. Nagamochi. Robust cost colorings. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2008.
14. D. Gijswijt, V. Jost, and M. Queyranne. Clique partitioning of interval graphs with submodular costs on the cliques. *RAIRO Operations Research*, 41:275–287, 2007.
15. D. Golovin, A. Gupta, A. Kumar, and K. Tangwongsan. All-norms and all- L_p -norms approximation algorithms. Technical Report CMU-CS-07-153, Carnegie-Mellon University, 2007.
16. M. Grötschel, L. Lovász, and A. Schrijver. *Geometric algorithms and combinatorial optimization*, volume 2 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, second edition, 1993.
17. M. T. Hajiaghayi, M. Mahdian, and V. S. Mirrokni. The facility location problem with general cost functions. *Networks*, 42(1):42–47, 2003.
18. M. M. Halldórsson and J. Radhakrishnan. Greed is good: approximating independent sets in sparse and bounded-degree graphs. *Algorithmica*, 18:145–163, 1997.
19. E. Halperin and R. M. Karp. The minimum-entropy set cover problem. *Theoretical Computer Science*, 348(2-3):240–250, 2005.
20. D. S. Johnson. Approximation algorithms for combinatorial problems. *Journal of Computer and System Sciences*, 9:256–278, 1974.
21. S. Khot. Improved inapproximability results for maxclique, chromatic number and approximate graph coloring. In *Proc. Annual Symposium on Foundations of Computer Science (FOCS)*, pages 600–609, 2001.
22. H. A. Kierstead and A. V. Kostochka. A short proof of the Hajnal-Szemerédi theorem on equitable colouring. *Combinatorics, Probability and Computing*, to appear.
23. L. Lovász. On the ratio of optimal integral and fractional covers. *Discrete Mathematics*, 13:383–390, 1975.
24. C. Lund and M. Yannakakis. On the hardness of approximating minimization problems. *Journal of the ACM*, 41(5):960–981, 1994.
25. K. Munagala, S. Babu, R. Motwani, and J. Widom. The pipelined set cover problem. In *10th International Conference on Database Theory (ICDT)*, volume 3363 of *Lecture Notes in Computer Science*, pages 83–98. Springer-Verlag, 2004.
26. D. Nakamura and A. Tamura. A revision of Minty’s algorithm for finding a maximum weight stable set of a claw-free graph. *J. Oper. Res. Soc. Japan*, 44(2):194–204, 2001.
27. R. Raz and M. Safra. A sub-constant error-probability low-degree test, and a sub-constant error-probability PCP characterization of NP. In *Proc. Annual ACM Symposium on Theory of computing (STOC)*, pages 475–484, 1997.

A Proof of theorem 5 and corollary 2

The proof is similar to that of lemma 1. We consider the approximate MaxIS algorithm in which a ρ -approximate maximum independent set is chosen at each step. We consider a class C_i in an optimal coloring, of size c_i . The first vertex v_1 of C_i that is colored by the approximate MaxIS algorithm will be assigned a value a'_{v_1} at least c_i/ρ , since there exists an independent set of size c_i in the current graph. By iterating this argument, we obtain that $\sum_{v \in C_i} (a'_v)^p \geq \frac{1}{\rho^p} \sum_{j=1}^{c_i} j^p$. In the optimal coloring, the value of this color class is c_i^{p+1} . Hence the ratio is at most

$$\left(\frac{n^{p+1}}{\frac{1}{\rho^p} \sum_{j=1}^n j^p} \right)^{\frac{1}{p}} = \rho \left(\frac{n^{p+1}}{\sum_{j=1}^n j^p} \right)^{\frac{1}{p}}. \quad (20)$$

For positive values of p , combining with lemma 2 yields an approximation factor of $\rho(p+1)^{\frac{1}{p}}$.

We now prove the corollary for the minimum entropy set cover problem. Using a greedy algorithm for the maximum independent set, we have $\rho = (\Delta + 2)/3$ [18]. This ratio is valid for each step of the

algorithm, as the maximum degree of the graph cannot increase. From (2), the error term for the minimum entropy problem is at most

$$\lim_{p \rightarrow 0} \log_2 \left(\frac{\Delta + 2}{3} (p + 1)^{\frac{1}{p}} \right) = \log_2(\Delta + 2) + \log_2(e) - \log_2(3) < \log_2(\Delta + 2) - 0.14226.$$